


85724

17P.



(NASA-CR-180977) NETWORK EFFECTS FOR
REAL-TIME APPLICATIONS (NASA) 17 p Avail:
NTIS HC A02/MF A01 CSCL 09B

N87-26568

Unclas

G3/61 0085724

RIACS

Research Institute for Advanced Computer Science

Network Protocols for Real-Time Applications

Marjory J. Johnson

May, 1987

Research Institute for Advanced Computer Science
NASA Ames Research Center

RIACS TR 87.15

Network Protocols for Real-Time Applications

Marjory J. Johnson

May, 1987

Research Institute for Advanced Computer Science
NASA Ames Research Center

The Fiber Distributed Data Interface (FDDI) and the SAE AE-9B High Speed Ring Bus (HSRB) are emerging standards for high-performance token ring local area networks. FDDI was designed to be a general-purpose high-performance network. HSRB was designed specifically for military real-time applications. A workshop was conducted at NASA Ames Research Center in January, 1987 to compare and contrast these protocols with respect to their ability to support real-time applications. This report summarizes workshop presentations and includes an independent comparison of the two protocols. A conclusion reached at the workshop was that current protocols for the upper layers of the Open Systems Interconnection (OSI) network model are inadequate for real-time applications.

This work was supported by the National Aeronautics and Space Administration under Cooperative Agreement NCC 2-387 to the Universities Space Research Association (USRA).

Network Protocols for Real-Time Applications

*Marjory J. Johnson
Research Institute for Advanced Computer Science
NASA Ames Research Center
Moffett Field, California 94035*

1. Introduction

On January 22-23, 1987, RIACS and the Intelligent Systems Architecture Branch of the Information Sciences Division at NASA Ames Research Center co-sponsored a workshop to compare and contrast the Fiber Distributed Data Interface (FDDI) and SAE AE-9B High Speed Ring Bus (HSRB) token ring protocols, to discuss the design philosophies of these protocols, and to determine how end-user applications dictate communication requirements for local area networks. The emphasis of the workshop was on real-time applications. Since both FDDI and HSRB are emerging standards for fiber optic high-performance networks, NASA and various agencies within the Department of Defense are investigating these protocols to determine their suitability for future networking applications, such as the Space Station. There were 45 attendees at the workshop, representing NASA, DoD, defense contractors, vendors, and independent consultants.

The meeting consisted of presentations by end users, protocol designers, and network analysts. Discussion at the end of the meeting focused on identifying problems of supporting real-time and high-performance applications in a networking environment.

2. Presentations

There were ten presentations, grouped into three categories. The first three presentations were discussions of end-user requirements, the next four were discussions of the design philosophies of FDDI and HSRB, and the last three were analyses of network performance.

2.1. Marc Cohn, Northrop Corporation

How Avionics Requirements Drive Communications Network Standards

Today's avionics systems are bus systems with centralized control. Such systems are inadequate to meet today's needs, because they are able to support only a limited number of stations (approximately 16). Hence, it is desirable to replace these centralized-control bus systems with distributed systems, but the level of fault-tolerance and responsiveness must remain the same. Northrop Corporation is monitoring various network protocol standards efforts, including FDDI and HSRB, to determine what type of network system should be adopted for avionics use and to influence development of the standards.

Cohn identified the following requirements for avionics networks. The network would be relatively small (up to 100 stations), consisting of a single subsystem without bridges or gateways. At most two priority levels would be required initially. The aircraft environment is harsh, due to extreme temperature

variation, vibration and shock, and possible exposure to nuclear radiation. Physical components of an avionics network would have to be developed to withstand these hardships. Limitations on power supply, size, weight, volume, and space for cable routing cause additional problems that must be addressed. Functional requirements include high system availability, minimization of mean time to repair (thus requiring the ability to connect and disconnect stations from the network easily), interoperability, use of standards, ability to incorporate anticipated advancements in a cost-effective manner, flexibility to support varying applications needs (e.g., sensors, processors, display stations), provisions for redundancy, high reliability, fault detection and isolation capabilities, high performance, and a distributed management and control scheme.

2.2. Art Justice, Naval Ocean Systems Center

Militarizing Commercial LAN Standards

Justice discussed naval shipboard requirements. The Navy Shipboard Advanced Fiber Optic Embedded Network (SAFENET) Committee is investigating local area networks for shipboard use. Current shipboard systems are point-to-point, with multiple point-to-point links for redundancy. This design is no longer satisfactory, because the number of ports is insufficient to provide the number of connections that are needed. Hence, the Navy is investigating replacing their point-to-point shipboard systems with local area networks. Justice envisions that many subnetworks will be required, connected via a backbone.

Members of the SAFENET committee are participating in various standards-development efforts, to attempt to influence the design so that it will satisfy Navy requirements. There will probably be a mix of commercial and military equipment on board a ship, so it is desirable to use commercial protocols whenever possible. Major differences between commercial and military networks were identified as the degrees of survivability, reconfigurability, and robustness that are required. According to Justice, there are lots of good protocols; no one protocol is best. He identified the ability to interconnect equipment easily as being more important for shipboard networks than high performance.

The SAFENET committee also plans to investigate protocols for the upper layers of the Open Systems Interconnection (OSI) model. They feel that the OSI model was developed for the unreliable switching networks of the telecommunications world, and that it is not suitable for Navy applications.

2.3. Mike Kearney, NASA Johnson Space Center

Mission Control LAN Upgrade

NASA Johnson Space Center plans to upgrade computer facilities at Mission Control by installing local area networks to replace the mainframe architecture which has been used since construction of Mission Control during the Apollo program in the 1960's. Two networks will be installed, a real-time network and a general-purpose network, both using FDDI as the media access control protocol. Kearney described the proposed structure of the two networks. The general-purpose LAN will be installed in the 1988-1989 time frame; the real-time LAN will follow in the 1990-1991 time frame. The real-time LAN will handle telemetry data. The general-purpose LAN will support all other functions. Transmission of time-critical command data is included on the general-purpose LAN, because it is felt that the real-time network will be too overloaded

handling telemetry data to be able to insure timely delivery of the command data without having to utilize priorities.

2.4. Manvel Geyer, Westinghouse

SAE AE-9B High Speed Ring Bus - Background

Geyer described the background and motivation for developing the SAE AE-9B High Speed Ring Bus, which was designed to replace MIL-STD 1553. Two subgroups of the Society of Automotive Engineers (SAE) have collaborated in this effort. The HART (High Speed Data Bus Applications and Requirements Task) subgroup produced a set of requirements for military real-time applications; the TAP (Topology and Protocol) subgroup determined that only the token ring topology could satisfy the HART requirements and that no existing protocol would satisfy all of them. Hence, the SAE AE-9B committee is developing the HSRB protocol, utilizing features of the IEEE 802.5 and the FDDI token ring protocols, both of which satisfied some, but not all, of the HART requirements. Geyer identified two major differences between HSRB and FDDI: the number of bits of internal station delay (4 to 6 with HSRB vs. up to 60 with FDDI) and type of access protocol (prioritized for HSRB vs. bandwidth sharing for FDDI). A more complete discussion of similarities and differences between the two protocols appears in sections 3.1 and 3.2.

2.5. Brian Kroeger, Westinghouse

SAE AE-9B High Speed Ring Bus - Tutorial

Kroeger highlighted the functional aspects of the HSRB protocol. He began by giving a general overview of the HART requirements for military real-time applications. These requirements include deterministic low message latency*, high throughput efficiency**, fault tolerance, and distributed control. HSRB was designed specifically to meet these performance requirements. In particular, HSRB was designed to operate on a bit-by-bit basis to reduce internal station delay and hence to decrease message latency.

Kroeger presented formulas for determining maximum latency for HSRB for a message of a given priority. It is possible to compute this time because messages are serviced in a strict priority order*** and because each station is allowed to transmit only a single message each time it captures the token.

Westinghouse has implemented a three-station prototype HSRB system. They are currently designing an HSRB controller chip, with delivery targeted for first quarter 1988.

*A message (referred to as a frame in the FDDI standards documents) is the unit of transmission over the channel. Message latency is the time from when a bus interface unit is notified of a message to send until the time the receiving bus interface unit is finished copying the message.

**Throughput efficiency is defined as the ratio of the amount of transmitted user-defined data to the bus data rate.

***Priorities are ignored immediately following transmission of a short message. Hence, strict priority order might not be followed if the short message option is selected.

2.6. Jim Hamstra, StanaTek

FDDI Design Tradeoffs

Hamstra is one of the primary developers of the FDDI protocol. He presented his perspective of the design rationale for FDDI. The design goals he identified, in order of importance, are utility, interoperability (among different FDDI implementations, with other LANs, and with standard upper-layer protocols), flexibility (of both applications that can be supported and physical configuration of the network), integrity, performance, extensibility, and agreeability. He emphasized the broad nature of the goals and the broad industry support in the development of the standard. Hamstra stressed that performance appears near the end of the list of design goals. Although FDDI is clearly a high-performance protocol, performance was not the most important design goal for FDDI, in direct contrast to the design philosophy for HSRB.

Hamstra discussed FDDI design decisions in detail, identifying the design goals supported by each decision. Hardware issues were addressed, as well as protocol issues.

2.7. Larry Green, Advanced Computer Communications

Impact of the Fiber Distributed Data Interface

Green is a member of the ANSI X3T9.5 committee that is developing the FDDI standard. His presentation emphasized the broad participation of vendors and end-users in the development of FDDI standards and products and in the wide range of possible applications for which FDDI would be suitable.

Green presented an overview of the driving forces behind the development of FDDI, including the demand for local area network technology beyond the limits of the IEEE 802 suite of protocols (in terms of performance, size, reliability, and functionality), the current trend toward standards preceding rather than following product development, the widespread participation of vendors in the development of FDDI, the possibility of obtaining high performance at low cost due to technological advances, and the wide variety of possible applications for which FDDI is suitable.

Green presented a lengthy list of vendors who are developing and will be marketing FDDI products. It is possible to purchase an FDDI prototype today, although the cost is prohibitively high. Green projected that the cost of FDDI will decrease just as rapidly as the cost of Ethernet did after its development, because FDDI enjoys the same widespread support and has the same prospects of revolutionizing the industry as Ethernet did.

2.8. Ken LaBel, NASA Goddard Space Flight Center

Idiosyncrasies of a Real Star

Sperry Flight Systems, Phoenix, Arizona, has developed a media access control protocol for a 100 megabit/second fiber optic star topology network, called the Star*Bus under contract from NASA Goddard Space Flight Center. The media access control protocol is a carrier sense multiple access/collision detection/collision resolution via time-slot (CSMA/CD/TS) protocol. Sperry delivered a three-node Star*Bus network to Goddard in September, 1985; Goddard has since been evaluating the system. LaBel's presentation described problems caused by limitations of buffer space within each station. Such problems

illustrate how actual results can differ considerably from expected theoretical results, regardless of the protocol under consideration.

Station architecture for the Star*Bus provides a single receive buffer in the station. A received message will occupy this buffer until it can be processed and absorbed into the host. This absorption time depends primarily on the length of the message. During this absorption time, any other message that arrives at this station is rejected. A rejected message will be retransmitted a specified number of times before it is discarded by the source as being undeliverable. It is easy to construct scenarios in which one or more stations will hog the channel, while others are starved. LaBel presented tables comparing the percentage of station access for individual stations to message size. For large message sizes, almost complete starvation is experienced by some of the stations. As message size is decreased, this starvation phenomenon is alleviated. This is because the probability that the single receive buffer is emptied before a new message arrives at the station increases as packet size decreases. In an experiment conducted with two stations transmitting to a single receiving station, the transmitting stations received equal access to the channel when the message size was 500 bytes or less.

2.9. Christina Berggren, IBM

Communications for Distributed Real-Time Fault-Tolerant Systems

Berggren is a member of the SAE AE-9B committee. She presented results of her analysis of performance of both the HSRB and the FDDI protocols.

Berggren's protocol analysis has been directed primarily towards determining how well various protocols (including upper-layer protocols) satisfy the latency requirements presented in the HART document. During her presentation, Berggren established a requirement for guaranteed latency of no more than 5 ms for an end-to-end communication through the network. According to Berggren, standards for real-time total communications do not exist at this time, i.e., current upper-layer protocols are inadequate. For this reason IBM is developing a three-layer communications architecture to support real-time applications, in conjunction with the HSRB effort.

Berggren identified the driver for LAN access as being a 16 ms sampling interval, with allowable jitter of approximately half a millisecond, to distribute a timing signal around the ring. This tight control on jitter appeared to be a major reason that the SAE AE-9B committee feels that FDDI's timed token protocol does not satisfy the HART requirements.

Berggren compared the FDDI and HSRB protocols on the basis of message latency. She presented results from some benchmark studies. An interpretation of the results she obtained is presented in section 3.5.

2.10. Kevin Mills, National Bureau of Standards

NBS Protocol Performance Program

NBS has conducted some network performance studies, to determine the appropriateness of OSI protocols for support of real-time applications. A few years ago Manufacturing Automation Protocol (MAP) developers expressed concern that existing OSI protocols wouldn't meet their real-time performance requirements. At that time, there was no data available to support or refute their concerns. NBS was asked to obtain some data.

NBS conducted experiments to measure end-to-end response on a testbed network. See [11] for a description of the work. The protocol suite implemented on the testbed consisted of CSMA/CD at the media access control layer, LLC 1 at the link layer, Transport Class 4 at the transport layer, with two applications at the top. The network, session, and presentation layers were all null. Measurements taken on the testbed showed that response time was inadequate to support real-time applications. The best one-way delay measured was 18 ms. Mills remarked that it wasn't just a matter of implementation of the protocols, but that the protocols themselves were unsatisfactory. Each layer contains some quality of service parameters, but there is no way to enforce quality of service except at the lowest layer, i.e., the media access control layer.

Mills identified the following important issues involving the use of OSI protocols for real-time applications: technology limitations, quality of service issues (e.g., how can the requested quality of service be guaranteed), protocol layering issues (e.g., how many layers are needed, should they really be implemented as layers), network management issues (with respect to both fault tolerance and performance), and tomorrow's technology issues (e.g., is the OSI model suitable to support high-performance protocols like FDDI and HSRB, or is something entirely different needed).

3. Comparison between FDDI and HSRB

It is assumed that the reader is familiar with the basic properties of the media access control protocols for FDDI, HSRB, and IEEE 802.5 [1,3,5,6,7,8,12,13,14]. HSRB combines the reservation/priority access scheme and the master clock synchronization scheme of 802.5 and the dual counter-rotating rings and 4B/5B encoding scheme of FDDI.

Major similarities and differences are discussed below.

3.1. Similarities

Both media access control protocols are high-performance protocols. FDDI is designed specifically for fiber optics; HSRB may use various physical media, including fiber optics. FDDI is a 100 megabit/second protocol; HSRB is presently intended to operate at 80 megabits/second with fiber optic media. Both FDDI and HSRB have been designed to accommodate higher transmission speeds without requiring any modification to the access protocols.

Both protocols provide guaranteed latency, FDDI via the timed token mechanism and HSRB via the reservation/priority scheme and the regulation that each station may transmit only a single message during each token-holding period.

Both protocols provide special attention to reliability issues, including providing redundancy and specifying reconfiguration techniques. Both physical configurations specify the use of dual counter-rotating rings.

3.2. Major Differences

HSRB was designed for rings that are limited in physical size to 128 stations and a 3 kilometer ring length. In contrast, FDDI can theoretically support an unlimited number of stations. Default parameter values for FDDI were calculated to support up to 1000 physical connections with a total fiber path length of 200 kilometers.

HSRB uses a reservation/priority access scheme, i.e., messages are serviced in order of priority.* FDDI is a timed token protocol; station timers cooperatively attempt to maintain a specified token rotation time by using the observed network load to regulate the amount of time that a station may transmit. Lower priority messages may be transmitted only if the overall network load is light enough to support it. FDDI supports message priorities, but messages among all the stations are not serviced in strict priority order.

The HSRB access protocol limits each station to sending a single message at each access opportunity. The FDDI protocol allows stations to transmit multiple messages, as long as station timer values permit.

Internal station delay is specified for HSRB as 4 to 6 bits, i.e., 40 to 60 nanoseconds at a 100 megabit per second transmission rate. For FDDI the delay is a maximum of 60 bits or 600 nanoseconds. HSRB requires more expensive hardware than FDDI in order to achieve such a small internal station delay. The two protocols use different synchronization schemes also. HSRB uses a master clock synchronization scheme; a synchronization adjustment is made in only one station, the station containing the master clock. FDDI uses an independent link synchronization scheme; each station contains an elasticity buffer, so that a synchronization adjustment is made within each station. Iyer and Joshi [9,10] and Salwen [15] have made some interesting observations about hardware issues and synchronization schemes. It should be noted that master clock synchronization is not feasible for large rings.

FDDI is designed to interface with standard higher-layer protocols, such as the IEEE 802.2 link layer protocol. HSRB is not. IBM is developing its own three layer network architecture, which includes HSRB as the bottom layer.

The FDDI media access control (MAC) document [8] is now an ANSI standard and has been forwarded to ISO to become an ISO standard. Other parts of the standard are nearing completion. Only the Station Management Document [7] is still in the formative stage. It is possible to purchase an FDDI prototype today, but it is prohibitively expensive. In contrast, although Westinghouse has developed a three-node in-house testbed, the HSRB protocol is still undergoing design changes.

3.3. Different Design Philosophies

The HSRB media access control protocol was designed to satisfy the HART requirements for military real-time applications. In contrast FDDI was designed by the commercial community to be a general-purpose high-performance protocol. There was general agreement at the workshop that military and/or space environments impose special requirements on a local area network. Designers of FDDI felt that if the PMD (physical media dependent) portion of the draft standard [6] were replaced with a ruggedized PMD, then FDDI would be suitable for military applications. That is, the access protocol would remain unchanged, but the physical part of the network would be adapted to the more strenuous military environment.

There was general consensus at the workshop that a design toward a

*Strict priority order might not be followed if the short message option is selected.

specific goal will always satisfy that goal better than a general-purpose design. However, as Hamstra discussed during his presentation, there are many design goals other than performance which may be worth considering. For example, for Space Station applications, reliability, flexibility, and extensibility are all more important than high performance. As another example, Cohn and Justice stressed the importance of interoperability for their applications.

3.4. Pros and Cons of Timed Token Protocols

A criticism of FDDI that was expressed during the workshop was the necessity of assigning values to the timing parameters. Some attendees considered this to be a serious obstacle to their use of FDDI. Hamstra responded that timing parameters have to be computed somewhere and enforced somewhere, regardless of the particular media access control protocol that is used. In HSRB, the computation must be done by the system designer. Enforcement of an upper bound on message latency would be difficult with HSRB, because message latency is dependent on message length and on the number of stations with highest priority messages to transmit. As Hamstra noted, simple formulas are given in the FDDI MAC document that can be used to determine appropriate settings for timing parameters. Once this is done, the MAC protocol enforces an upper bound on message latency via the timed token mechanism.

An additional consideration is the flexibility of the timed token approach. With FDDI a station may send multiple messages during a single access opportunity if time permits, while with HSRB each station is allowed to transmit only a single message. Hence, FDDI will service bursty traffic more efficiently. In addition, through unequal synchronous bandwidth allocations with FDDI, different station requirements for transmission of highest priority messages can easily be accommodated.

3.5. Limitations of Existing Performance Comparison Studies

Some performance comparisons between FDDI and HSRB have been conducted by Westinghouse [4] and IBM (reported on by Berggren at the workshop) under the auspices of the SAE AE-9B committee, to justify the development of a new protocol to satisfy the HART requirements. These studies have consisted of comparing message latency of the two protocols for a limited number of static traffic scenarios. Some of the scenarios that were selected were based on unrealistic assumptions regarding distribution of priorities (e.g., only one station on the ring has a highest priority message to send at a given time), message generation (e.g., messages are generated at constant time intervals and at exactly the same time at all stations), and queueing delays (e.g., queueing delays are nonexistent because all messages on the ring are transmitted before the next messages are generated). In addition, priority threshold settings for FDDI were chosen inappropriately in some of these scenarios, so as to block lowest priority traffic from ever being transmitted. Meaningful conclusions cannot be reached from such limited studies.

To obtain a valid comparison of the two protocols, they must be tested in a variety of situations. Since both protocols are still under development, the best means of comparison at this time is simulation. Simulations of FDDI are available at NASA Ames Research Center, IBM - Zurich, and Synetics (through the Naval Ocean Systems Center). Synetics is planning to model HSRB also. A

comprehensive comparison of the two protocols should address issues such as reliability, flexibility, interoperability, extensibility, and cost, as well as performance, since these other properties might be more important than performance in certain applications.

It remained unclear at the workshop whether or not FDDI would meet the HART requirements. Since the workshop, Cohn [2] has addressed this question and has concluded that FDDI does indeed meet the HART requirements.

3.6. Sample Scenarios

In existing performance comparisons, the shorter message latency for HSRB over FDDI can be attributed primarily to the significantly smaller internal station delay for HSRB. However, it is easy to define scenarios where FDDI will provide shorter message latency than HSRB, because of FDDI's flexibility. Sample scenarios are presented below. HSRB provides superior performance in the first scenario; FDDI provides superior performance in the second scenario; each protocol provides superior performance in different ways in the last scenario. Hence, results obtained by considering isolated scenarios cannot be used to establish the general superiority of one protocol over the other.

The following general assumptions apply to all the scenarios presented below: all messages are long enough to fill up the ring, all messages are the same size, there is no other traffic on the ring, the transmission rate is 100 megabits per second, and the header size and token size are the same for both FDDI and HSRB.*

Scenario 1

- 600 meter ring
- 8 active stations, equally spaced around the ring
- each station has one message to transmit
- all messages have the same priority

Set the target token rotation time parameter for FDDI so that all messages are transmitted during a single token rotation. All messages will be serviced during a single token rotation with HSRB also. The difference between internal station delay between the two protocols clearly means that all messages are transmitted more quickly with HSRB than with FDDI.

Scenario 2

- 600 meter ring
- 8 active stations, equally spaced around the ring
- each station has three messages to transmit
- all messages have the same priority

Set the target token rotation time parameter for FDDI so that all messages are transmitted during a single token rotation. Three token rotations will be

*The assumption that transmission rate, header size, and token size are the same for FDDI and HSRB enables a more direct comparison of the two access schemes.

required for the HSRB, since each station may transmit only one message during each access opportunity. An easy computation establishes that all messages are transmitted more quickly with FDDI than with HSRB, because fewer token rotations are required. Hence, FDDI provides superior performance in this situation.

Scenario 3

- 600 meter ring
- 8 active stations, equally spaced around the ring
- each station has one message to transmit
- each message has a different priority
- the messages are located in descending priority order in one direction around the ring, while the token rotates in the opposite direction

Set the target token rotation time and priority threshold parameters for FDDI so that all messages are transmitted during a single token rotation. Since the messages are located so that their priorities are backwards in relation to the direction of token rotation, the token will have to rotate almost all the way around the ring after each message is transmitted in order to reach the station with the next highest priority message to transmit. An easy computation establishes that all messages are transmitted more quickly with FDDI than with HSRB, due to the larger number of token rotations required for HSRB. However, the higher priority messages will be transmitted more quickly with HSRB than with FDDI, because the messages will be serviced in priority order.

4. OSI Network Architectural Issues

Both FDDI and HSRB are protocols for the bottom one and a half layers of the OSI network model. Standards exist and/or are being developed for protocols for the higher layers of the OSI model. FDDI is being designed to interface with these standard higher-layer protocols; HSRB is not. Berggren explained that IBM is designing its own three-layer network architecture (based on the OSI model) to use in conjunction with HSRB.

It was generally agreed at the workshop that the difference in message latency for the two protocols as indicated in Berggren's benchmark studies is insignificant. The processing delays introduced by the upper layers of the OSI network model will dwarf differences in latency between FDDI and HSRB at the MAC layer. The question was raised as to whether the OSI model can support real-time applications. It was stressed that the problem is not just one of implementation; it is also a problem with the architecture itself. Higher layer protocols already form the bottleneck for networks using current technology. Hence, increasing the transmission rate with FDDI or HSRB networks does nothing to improve overall network performance.

We discussed possible approaches by which our workshop group might develop a suitable architecture and protocols to support real-time applications. One approach would be to follow the MAP ideology, adhering to the OSI model, but allowing null layers and reduced functionality in other layers. Another approach would be to start from scratch and develop a set of upper-layer network protocols specifically for real-time needs.

5. Conclusions and Future Efforts

The workshop provided a forum for people with common interests in high-performance and real-time networking to discuss networking needs and requirements. As a result of our discussions, we identified the important issues to address in providing support for real-time applications. We did not obtain solutions, but we did pinpoint the problems.

The first set of presentations at the workshop provided insight into end-user requirements. Reliability, interoperability, flexibility, and extensibility were identified as being just as important as high performance. There was also general acceptance of the importance of management to the successful operation of a network, and an appreciation of the difficulties in agreeing to management concepts.

The next set of presentations provided insight into the motivations behind the development of each protocol. HSRB was designed specifically for real-time military applications. FDDI is a general-purpose high-performance protocol, but it is still a viable candidate for real-time applications.

The third set of presentations addressed performance issues. A comprehensive comparison of FDDI and HSRB has not yet been done; an objective simulation study would be useful. While discussing the significance of the differences in message latency provided by the two protocols, we realized that the real bottleneck in networking lies in the structure of the OSI model and in existing upper-layer protocols.

Several of the speakers at the workshop stressed the desirability of using networking standards. NASA and DoD are relying heavily on the commercial world to provide solutions to their networking problems. The need for NASA and DoD participation in standards development was recognized by both government, defense contractor, and commercial representatives. This is the only way that NASA and DoD can influence the development of the standards and the only way that protocol developers can find out what NASA and DoD requirements really are.

The government lags behind the commercial world in computer networking technology. Several of the DoD representatives expressed the recognition that there is currently a window of opportunity during which to impact networking for the next 10 to 20 years within NASA and DoD. Current opportunities include network design for the Space Station, redesign of shipboard systems for the Navy, and redesign of Mission Control at NASA JSC.

A strong commitment was expressed within the group to continue joint discussions to determine a suitable network architecture to support real-time applications. A follow-on workshop (with limited attendance) is planned at the National Bureau of Standards in June. This future workshop will focus on identifying the requirements of real-time systems, on the applicability of the OSI model to real-time systems, and on defining a suitable model for real-time communications.

6. References

1. ANSI/IEEE 802.5, "Token Ring Access Method and Physical Layer Specifications."
2. M. Cohn, "The Conformance of the ANSI FDDI Standard to the SAE-9B HART High Speed Data Bus Requirements for Real-Time Local Area Networks," May 8, 1987, accepted for presentation at the November, 1987 meeting of the SAE-9 committee.
3. R. J. DelCoco, B. W. Kroeger, J. J. Kurtz, "An Overview of SAE AE-9B High Speed Ring Bus (HSRB) Performance."
4. R. J. DelCoco and J. J. Kurtz, "Comparison of the ANSI FDDI and SAE AE-9B Token Ring LAN's," preliminary draft.
5. "FDDI Physical Layer Protocol," Draft Proposed American National Standard, X3T9.5/83-15, Rev. 13, August 22, 1986.
6. "FDDI Physical Medium Dependent," Draft Proposed American National Standard, X3T9.5/84-48, Rev. 7, February 20, 1987.
7. "FDDI Station Management," Draft Proposed American National Standard, X3T9.5/84-49, Rev. 2.1, April 16, 1987.
8. "FDDI Token Ring Media Access Control," Draft Proposed American National Standard, X3T9.5/83-16, Rev. 10, February 28, 1986.
9. V. Iyer and S. P. Joshi, "FDDI's 100 M-bps Protocol Improves on 802.5 Spec's 4-M-bps Limit," EDN, pp. 151-160, May 2, 1985.
10. V. Iyer and S. P. Joshi, "Streamlining Protocols at High Data Rates," Proceedings of Wescon/84, Session 2, Paper 4.
11. K. Mills, M. Wheatley, and S. Heatley, "Real-Time Performance Predictions: Transport Over LANs," Proceedings of Advanced Seminar on Real-Time Local Area Networks, Bandol, France, April, 1986, pp. 345-358.
12. F. E. Ross, "FDDI - A Tutorial," IEEE Communications Magazine, Vol. 24, No. 5, pp. 10-17, 1986.
13. F. E. Ross and R. K. Moulton, "FDDI Overview - a 100 Megabit per Second Solution," Proceedings of Wescon/84, Session 2, Paper 1.
14. "SAE Draft High Speed Ring Bus (HSRB) Standard," October 27, 1986.

15. H. Salwen, "Phase Jitter in Token Passing Rings," Proceedings of 11th Conference on Local Computer Networks, October, 1986, pp. 11-15.

7. Workshop Attendees

Richard Bailly, Vitro Laboratories
Christina Berggren, IBM
Larry Bergman, JPL
Frank Bruns, Proteon
Peter Carrillo, McDonnell Douglas Corporation
Tammy Chan, Honeywell
Greg Chesson, Silicon Graphics, Inc.
Marc Cohn, Northrop Corporation
Jim Comeau, Hughes Aircraft Co.
John DeRuiter, Sperry Flight Systems
Barry Driggs, IBM
Uldis Duselis, Naval Avionics Center
Farzin Firoozmand, Advanced Micro Devices
Manvel Geyer, Westinghouse DEC
Andrew Glista, Naval Air Systems Command
Harry Gold, FCDSSASD
Terry Grant, NASA Ames Research Center
Larry Green, Advanced Computer Communications
Bob Grow, GP Associates
Jim Hamstra, StanaTek
Tom Hipp, Honeywell
Ed Howard, Naval Ocean Systems Center
Marjory Johnson, RIACS
Art Justice, Naval Ocean Systems Center
Mike Kearney, NASA Johnson Space Center
Robert Kolacki, Space & Naval Warfare Systems Command
Brian Kroeger, Westinghouse DEC
Ken LaBel, NASA Goddard Space Flight Center
Hoyt Layson, Honeywell
Eugene Lew, Sperry Flight Systems
John Love, UNISYS
Kevin Mills, National Bureau of Standards
Nicholas Murray, NASA Langley Research Center
Louis Neri, Naval Air Engineering Center
Dave Pfahler, Hughes Aircraft Company
Rama Raghavender, TRW
Bill Reed, NASA Johnson Space Center
Ken Rehm, IBM
Peter Relan, UNISYS
Andy Schreckenghost, Ford Aerospace
Richard Steinberger, Martin Marietta Aerospace
Richard Swaney, Hughes Aircraft Company
Alfred Weaver, University of Virginia
Pat Wilhelm, McDonnell Douglas Corporation
Don Zurek, Sperry Corporation